arXiv:1401.2920v1 [cs.DC] 13 Jan 2014

The EFTOS Voting Farm: A Software Tool for Fault Masking in Message Passing Parallel Environments

Vincenzo De Florio, Geert Deconinck, Rudy Lauwereins

Katholieke Universiteit Leuven
Electrical Engineering Dept. – ACCA
Kard. Mercierlaan 94 – B-3001 Heverlee – Belgium

Abstract. We present a set of C functions implementing a distributed software voting mechanism for EPX or similar message passing environments, and we place it within the EFTOS framework (Embedded Fault-Tolerant Supercomputing, ESPRIT-IV Project 21012) of software tools for enhancing the dependability of a user application. The described mechanism can be used for instance to implement restoring organs i.e., N-modular redundancy systems with N-replicated voters. We show that, besides structural design goals like fault transparency, this tool achieves replication transparency, a high degree of flexibility and ease-of-use, and good performance.

Keywords: Software Fault Masking, Fault Tolerance, Voting Techniques, High-Performance Computing.

1 Introduction

A well-known approach to achieve fault masking and therefore to hide the occurrence of faults is the N-modular redundancy (NMR) technique (see for instance [9]), valid both on hardware and at software level. To overcome the shortcoming of having one voter, whose failure brings to the failure of the whole system even when each and every other module is still running correctly, it is possible to use N replicas of the voter and to provide N copies of the inputs to each replica, as described in Fig.1. This approach exhibits among others the following properties:

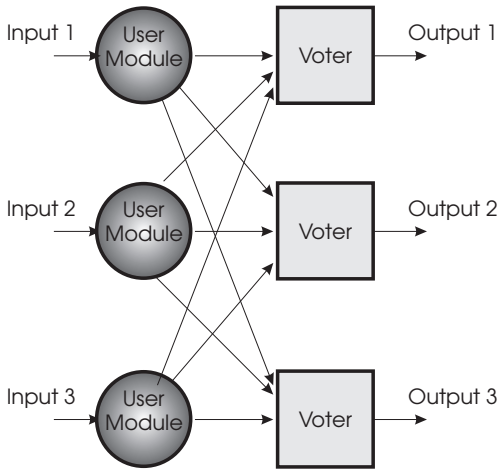

Fig. 1. A "restoring organ" [9] i.e., a N-modular redundant system with N voters, when $N = 3$.

1. Depending on the voting technique adopted in the voter, the occurrence of a limited number of faults in the inputs to the voters may be masked to the subsequent modules [12]; for instance, by using majority voting, up to $\llcorner (N-1)/2 \lrcorner$ faults can be made transparent.
2. If we consider a pipeline of such systems, then a failing voter in one stage of the pipeline can be simply regarded as a corrupted input for the next stage, where it will be restored.

The resulting system is easily recognizable to be more robust than plain *N*MR, for it does no more exhibit single points of failures. Dependability analysis confirms intuition. Property 2. in particular explains why such systems are also known as "restoring organs" [9].

From the point of view of software engineering, this system though has two major drawbacks:

- Each module in the *N*MR must be aware of and responsible for interacting with the whole set of voters;
- The complexity of these interactions, which is a function increasing quadratically with N, the cardinality of the voting farm, burdens each module in the *N*MR.

As a consequence, it firstly appeared difficult to us to design a software mechanism which, besides reaching design goals like fault transparency (i.e., fault masking) and efficiency, were also able to achieve replication transparency, ease of use, and flexibility.

In order to reach the full set of these requirements, we slightly modified the design of the system as described in Fig.2: Now each module only has to interact with, and be aware of *one* voter, regardless the value of N. Moreover, the complexity of such a task is fully shifted to the voter.

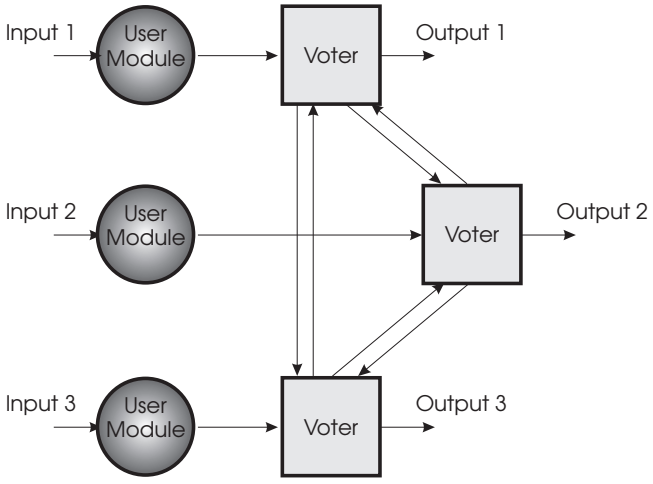

Fig. 2. Structure of the EFTOS voting farm mechanism for a *N*MR system with $N = 3$ (also known as TMR, or triple modular redundancy system.)

We adopted this approach during the design and development of the voting farm mechanism, a class of C functions which is part of the EFTOS framework (Embedded Fault-Tolerant Supercomputing, ESPRIT-IV Project 21012). In this paper we briefly draw a picture of EFTOS and place the voting farm into it; then we describe the design of the voting farm and show how such tool proved to fulfill the whole of

our design goals and requirements. We also describe how the user can exploit it to easily set up systems consisting of redundant modules and based on voters. A few notes on future developments and portability conclude this work.

2 EFTOS and its Framework

The overall object of the ESPRIT-IV Project 21012 EFTOS [5,6] (Embedded Fault-Tolerant Supercomputing) is to set up a software framework for integrating fault tolerance into embedded distributed high-performance applications in a flexible and easy way. The EFTOS framework currently runs on a Parsytec CC system [1], a distributed-memory MIMD supercomputer consisting of powerful processing nodes based on PowerPC 604 at 133MHz, dedicated high-speed links, I/O modules, and routers. As part of the Project, this framework is currently being ported to Microsoft Windows NT / Intel PentiumPro and TEX / DEC Alpha platforms so to fulfill the requirements of the EFTOS application partners. We herein constantly refer to the version running on the CC system.

The main characteristics of the CC system are the adoption of the thread processing model and of the message passing communication model: communicating threads exchange messages through a proprietary message passing library, called EPX [2] (for *Embedded Parallel extensions to uniX*). A noteworthy feature of the EPX environment, which revealed to be very useful for the voting farm, is the EPX so-called "initial loading mechanism," which spawns the same executable image of the user application on each processing node of the user partition (a sort of "parallel fork()" [8]). As a final note on EPX, we remark that it adopts the concept of "virtual links" to build point-to-point connections between arbitrary threads within the processor pool, and that of "local links" to create similar connections between threads running on a same node—the only noticeable difference being of course in terms of performance. Once a connection among any two threads has been set up, the involved threads refer to it by means of a link, and use it to send and receive messages along the same connection. Send()'s and Receive()'s are synchronous and blocking—this latter attribute being a potential source of problems from the viewpoint of fault tolerance. Receive()'s are "better," in the sense that it is possible to specify a time-out that, once reached, unblocks the caller regardless the operation has reached completion. Such a functionality is missing in the Send() function.

Through the adoption of the EFTOS framework, the target embedded parallel application is plugged into a hierarchical, layered system whose structure and basic components are:

- At the lowest level, a set of parametrisable functions managing error detection (Dtools) and error recovery (Rtools). A typical Dtool is a watchdog timer thread or a trap-handling mechanism; a Rtool is e.g., a fast-reboot thread capable of restarting a single node or a set of nodes. These basic components are plugged into the embedded application to make it more dependable. EFTOS supplies a number of these Dtools and Rtools, plus an API for incorporating user-defined EFTOS-compliant tools;
- At the middle level, a distributed application called *DIR net* (detection, isolation, and recovery network) [14] is available to coherently combine Dtools and Rtools, to ensure consistent strategies throughout the whole system, and to play the role of a backbone handling information to and from the fault tolerance elements;
- At the highest level, these elements are combined into dependable mechanisms i.e., methods to guarantee fault-tolerant communication, the voting farm mechanism, etc.

During the lifetime of the application, this framework guards it from a series of possible deviations from the expected activity; this is done either by executing detection, isolation, and reconfiguration tasks, or by means of fault masking—this latter being provided by the EFTOS voting farm, which we are going to describe in the Section to follow. As a last remark, the EFTOS framework appears to the user as a library of functions written in the C programming language.

3 The EFTOS Voting Farm

The basic component of our tool is the *voter* (see Fig.3) which we define as follows:

> A voter is a local software module connected to *one* user module and to a farm of fully interconnected fellows. Attribute "local" means that both user module and voter run on the same processing node.

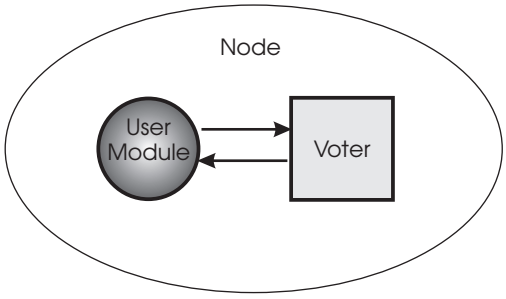

Fig. 3. A user module and its voter. The latter is the only member of the farm of which the user module should be aware of: messages will flow only between these two ends. This has been designed so to minimize the burden of the user module and to keep it free to continue undisturbed as much as possible.

As a consequence of the above definition, the user module has no other interlocutor than its voter, whose tasks are completely transparent to the user module. It is therefore possible to model the whole system as a simple client-server application [4]: on each user module the same client protocol applies (see §3.1) while the same server protocol is executed on every instance of the voter (see §3.2).

3.1 The Client-Side of the Voting Farm

Table 1 gives an example of the client-side protocol to be executed on each processing node of the system in which a user module runs: a well-defined, ordered list of actions has to take place so that the voting farm be coherently declared and defined (in the sense specified in [13]), described, activated, controlled, and queried. In particular, *describing* a farm stands for creating a static map of the allocation of its components; *activating* a farm substantially means spawning the local voter (§3.2 will shed more light on this); *controlling* a farm means requesting its service by means of control and data messages; finally, a voting farm can also be *queried* about its state, the current voted value, etc.

As already mentioned, the above steps have to be carried out in the same way on each user module: this coherency is transparently supported by the "initial load mechanism" of EPX [1].

This protocol is available to the user as a class-like collection of functions dealing with opaque objects referenced through pointers. A tight resemblance with the FILE

set of functions of the standard C language library [10] has been sought so to shorten as much as possible the user's learning time. The FILE paradigm shows also that, though C is certainly not the best language for object-oriented programming, its support for data and function hiding[1], coupled with good software practice can combine effectiveness, efficiency, and the elegance of object-orientation.

The EFTOS voting farm adopts these principles and its API and usage closely resemble those of FILE. It also benefits from the use of the CWEB system of structured documentation [11] which we found an extremely useful design tool [7].

1	VotingFarm_t *vf;	/* declaration */
2	vf ← VF_open(objcmp);	/* definition */
3	$\forall i \in \{1, \ldots, N\}$: VF_add(vf, node$_i$);	/* description */
4	VF_run(vf);	/* activation */
5	VF_control(vf, VF_input(obj, sizeof(obj_t)), VF_output(link), VF_algorithm(VFA_WEIGHTED_AVERAGE), VF_scaling_factor(1.0));	/* control */
6	do {} while (VF_error==NONE ∧ VF_get(vf)==VF_REFUSED);	/* query */
7	VF_close(vf);	/* deactivation */

Table 1. An example of usage of the voting farm: note the resemblance with the FILE standard set of C language functions. objcmp() is a user-supplied function for comparing any two obj_t objects—its role is explained in §3.2. Note also how four messages are sent to the local voter in Step **5**: the input to be voted, the virtual link representing the thread to whom the voted output has to be sent, the voting algorithm, and an optional argument pertaining the algorithm. As a final remark, Step **6** is needed because one can only terminate a voting farm when the broadcast of the input value is over; any attempt to do that sooner results in a VF_REFUSED message. The loop also checks whether a time-out has occurred during a VF_get(), in which case the global variable VF_error is set to a value different from NONE.

3.2 The Server-Side of the Voting Farm: the Voter

The local voter thread represents the server-side of the voting farm. After the set up of the static description of the farm (Table 1, Step 3) in the form of an ordered list of processing node identifiers (integer numbers greater than 0), the server-side of our application is launched by the user by means of the VF_run() function. This turns the static representation of a farm into an "alive" (running, according to [3]) object, the voter thread.

This latter connects to its user module via a local link and to the rest of the farm via virtual links. From then on, in absence of faults, it reacts to the arrival of the user messages as a finite-state automaton: in particular, the input messages arrival triggers a number of broadcasts among the voters—as shown in Fig.4—which are managed through the distributed algorithm described in Table 2. When faults occur and affect up to $M < N$ voters, no arrival for more than Δt time units is interpreted as the symptom of a fault. As a consequence, variable input_messages is incremented as if a message had arrived, and its faulty state is recorded. This way we can tolerate up to $M < N$ errors at the cost of $M\Delta t$ time units. Note that even

[1] In our opinion these concepts, available long before the conception of the C++ programming language [13], must have been a powerful conceptual inspirer for this latter.

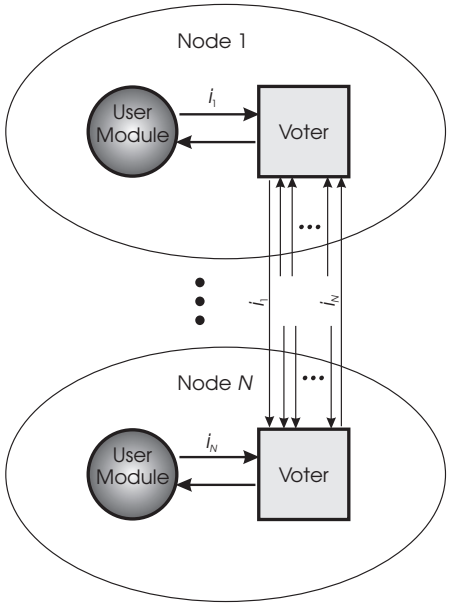

Fig. 4. The "local" input value has to be broadcasted to $N-1$ fellows, and $N-1$ "remote" input values have to be collected from each of the fellows. The voting algorithm takes place as soon as a full suite of values is available. Note that a distributed algorithm is needed to regulate at all times who has the right to broadcast and who has to receive.

though this algorithm tolerates up to $N-1$ faults, the voting algorithm may be able to cope with much less than that: for instance, majority voting fails in the presence of $\llcorner (N-1)/2 \lrcorner + 1$ or more faults. As another example, algorithms computing a weighted average of the input values consider all items whose "faulty bit" is set as zero-weight values, automatically discarding them from the average. This of course may also lead to imprecise results as the number of faults gets larger.

Besides the input value, which represents a request for voting, the user module may send its voter a number of other requests—some of these are used in Table 1, Step 5. In particular, the user can choose to adopt a voting algorithm out of the following ones:

- Formalized majority voter technique,
- Generalized median voter technique,
- Formalized plurality voter technique,
- Weighted averaging technique,

namely the voting techniques that were generalized in [12] to "arbitrary N-version systems with arbitrary output types using a metric space framework." To use these algorithms, a metric function can be supplied by the user when he/she "opens" the farm (Table 1, Step 2): this is exactly the same approach used in opaque C functions like e.g., bsearch() or qsort() [10]. A default metric function is also available.

Other requests include the setting of some system parameters and the removal of the voting farm (function VF_close()).

The voters' replies to the incoming requests are straightforward. In particular, a VF_DONE message is sent to the user module when a broadcast has been performed; for the sake of avoiding deadlocks, one can only control or close a farm after the VF_DONE message has been sent. Any failed attempt causes the voter to send a VF_REFUSED message. This is the rationale of Step 6 in Table 1.

Note how a function like VF_get() simply sets the caller in a waiting state from which it exits either on a message arrival or on the expiration of a time-out. Doing

```
1   voter_id ← who-am-i();        /* identify yourself (voter_id ∈ {1,...,N}) */
2   ∀i : valid_i ← TRUE;          /* all messages are supposed to be valid */
3   i ← input_messages ← 0;       /* keep track of the number of received input messages */
4   do {
5         Wait_Msg_With_Timeout(Δt);      /* wait for an incoming message or a timeout */
6         if ( Sender() == USER ) u ← i;  /* u points to the user module's input */
7         if ( ¬ Timeout ) msg_i ← Receive();/* read it */
8         else valid_i ← FALSE;           /* or invalidate its entry */
9         i ← input_messages ← input_messages + 1; /* count it */
10        if (voter_id == input_messages) Broadcast(msg_u);
11  } while (input_messages ≠ N);
```

Table 2. The distributed algorithm needed to regulate the right to broadcast among the N voters. Each voter waits for a message for a time which is at most Δt, then it assumes a fault affected either a user module or its voter. Function Broadcast() sends its argument to all voters whose id is different from voter_id. It is managed via a special sending thread to avoid the deadlock-prone `Send()`.

the other way around would have been more error prone because of the lack-of-timeout problem reported in Section 2.

4 Time and Resources Overheads of the Voting Farm.

All measurements have been performed running a restoring organ consisting of N processing nodes, $N = 1, \ldots, 4$ [5]. The executable file has been obtained with the `ancc` C compiler using the `-O` optimization flag. During the trials the CC system was fully dedicated to the execution of that application.

The application has been executed in four runs, each of which has been repeated fifty times, increasing the number of voters from 1 to 4. Wall-clock times have been collected. Averages and standard deviations are shown in Table 3.

number of nodes	average	standard deviation
1	0.000615	0.000006
2	0.001684	0.000022
3	0.002224	0.000035
4	0.003502	0.000144

Table 3. Time overhead of the voting farm for one to four node systems. The unit is seconds.

As of the overhead in resources, N threads have to be spawned, and N local links are needed for the communication between each user module and its local voter. The network of voters calls for another $\frac{N \times (N-1)}{2}$ virtual links.

5 Conclusions

A flexible, easy to use, efficient mechanism for software voting in message passing systems has been described. The tool, currently running on a Parsytec CC system, has been designed with portability in mind and is actually being ported to a PentiumPro/Windows NT and a Alpha/TEX platform. A special, "static" version is

being developed for this latter, which adopts the mailbox paradigm as opposed to message passing via virtual links.

We are currently considering some additional improvements and extensions of the voting farm, including the possibility for the user to supply voting algorithms of his/her choice, and a tighter link with the EFTOS fault tolerance backbone: in particular, the voting farm will inform the DIR net about the state of the voting sessions, namely who failed, and on which nodes this happened. This information shall be exploited by the "error diagnosis engine" [14] of the DIR net.

Acknowledgments. This project is partly sponsored by an FWO Krediet aan Navorsers, by the Esprit-IV Project 21012 EFTOS, and by COF/96/11. Vincenzo De Florio is on leave from Tecnopolis CSATA Novus Ortus. Geert Deconinck has a grant from the Flemish Institute for the Promotion of Scientific and Technological Research in Industry (IWT). Rudy Lauwereins is a Senior Research Associate of the Fund for Scientific Research - Flanders (Belgium).

References

1. Anonymous, "Parsytec CC Series—Cognitive Computing," Parsytec GmbH, Aachen, 1996.
2. Anonymous, Embedded Parix Programmer's Guide, *in* "Parsytec CC Series Hardware Documentation," Parsytec GmbH, Aachen, 1996.
3. N. J. Carriero and D. Gelernter, How to write parallel programs: a guide to the perplexed, *ACM Comp. Surv.* **21** (1989), 323–357
4. D. E. Comer and D. L. Stevens, "Internetworking with TCP/IP, Volume 3: Client-Server Programming and Applications," Prentice-Hall, Englewood-Cliffs, 1993.
5. G. Deconinck, V. De Florio, R. Lauwereins, W. Rosseel, and M. Truyens, The EFTOS Reference Guide and Cookbook, Deliverable 2.4.2, Project IT-21012 EFTOS, March 1997.
6. G. Deconinck, V. De Florio, R. Lauwereins, and T. Varvarigou, EFTOS: A Software Framework for More Dependable Embedded HPC Applications, accepted for presentation at the European Conf. in Parallel Processing (Euro-Par'97).
7. V. De Florio, The Voting Farm—A Distributed Class for Software Voting, Tech. Rep. Katholieke Universiteit Leuven ESAT/ACCA/1997/3, June 7, 1997.
8. K. Haviland and B. Salama, "UNIX System Programming," Addison-Wesley, Reading MA, 1987.
9. B. W. Johnson, "Design and Analysis of Fault-Tolerant Digital Systems," Addison-Wesley, New York, 1989.
10. B. W. Kernighan and D. M. Ritchie, "The C Programming Language (2nd ed.)," Prentice-Hall, Englewood Cliffs, 1988.
11. D. E. Knuth, "Literate Programming," Center for the Study of Language and Information, Leland Standard Junior University, 1992.
12. P. R. Lorczak, A. K. Caglayan, and D. E. Eckhardt, A Theoretical Investigation of Generalized Voters for Redundant Systems, *in* Proc. of the 19th Int. Sym. on Fault-Tolerant Computing (1989), 444–451.
13. B. Stroustrup, "The C++ Programming Language (2nd ed.)," Addison-Wesley, Reading MA, 1995.
14. M. Truyens, G. Deconinck, V. De Florio, W. Rosseel, and R. Lauwereins, The DIR net: an intelligent backbone for software fault-tolerance in embedded parallel applications, submitted to the 4th Int. Symp. on High-Performance Computer Architecture (HPCA-4), Las Vegas, Nevada, Feb. 1–4, 1998.

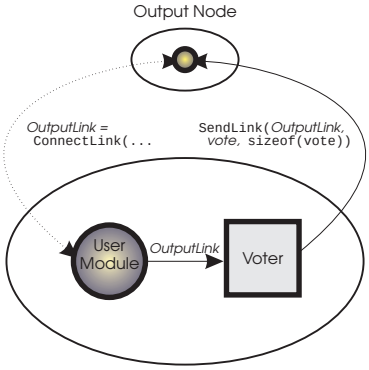
Output Node
OutputLink =
ConnectLink(...
SendLink(OutputLink,
vote, sizeof(vote))
User
Module
OutputLink
Voter